

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Beukema et al.	§	
	§	Group Art Unit: 2151
Serial No. 09/925,578	§	
	§	Examiner: Walsh, John B.
Filed: August 9, 2001	§	
	§	
For: Apparatus and Method for	§	
Implementing Multicast on a System	§	
Area Network Channel Adapter		

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

35525
PATENT TRADEMARK OFFICE
CUSTOMER NUMBER

APPEAL BRIEF (37 C.F.R. 41.37)

This brief is in furtherance of the Notice of Appeal, filed in this case on August 30, 2006.

A fee of \$500.00 is required for filing an Appeal Brief. Please charge this fee to IBM Corporation Deposit Account No. 09-0447. No additional fees are believed to be necessary. If, however, any additional fees are required, I authorize the Commissioner to charge these fees which may be required to IBM Corporation Deposit Account No. 09-0447. No extension of time is believed to be necessary. If, however, an extension of time is required, the extension is requested, and I authorize the Commissioner to charge any fees for this extension to IBM Corporation Deposit Account No. 09-0447.

REAL PARTY IN INTEREST

The real party in interest in this appeal is the following party: International Business Machines Corporation of Armonk, New York.

RELATED APPEALS AND INTERFERENCES

With respect to other appeals or interferences that will directly affect, or be directly affected by, or have a bearing on the Board's decision in the pending appeal, there are no such appeals or interferences.

STATUS OF CLAIMS

A. TOTAL NUMBER OF CLAIMS IN APPLICATION

Claims in the application are: 1-35

B. STATUS OF ALL THE CLAIMS IN APPLICATION

1. Claims canceled: none
2. Claims withdrawn from consideration but not canceled: none
3. Claims pending: 1-35
4. Claims allowed: none
5. Claims rejected: 1-35
6. Claims objected to: none

C. CLAIMS ON APPEAL

The claims on appeal are: 1-35

STATUS OF AMENDMENTS

An amendment after final was filed by Appellants on July 3, 2006, and was not entered by the Examiner according to an Advisory Action dated 8/2/2006.

SUMMARY OF CLAIMED SUBJECT MATTER

A. CLAIM 1 - INDEPENDENT

In a storage area network, a channel adapter receives a multi-cast packet and determines which local queue pairs are a part of the multi-cast group as identified by an identifier in the multi-cast data packet. The channel adapter replicates the data packet and delivers a copy of the data packet to each local queue pair that is part of the multi-cast group. Specifically, Claim 1 is directed to a method of multicasting a data packet in a system area network. A data packet is received, where the data packet includes an identifier of a multicast group. A plurality of queue pairs that are members of the multicast group are identified. The data packet is delivered to each of the plurality of queue pairs that are members of the multicast group (Specification page 9, line 14 – page 10, line 2; Figure 3, element 300; Specification page 11, line 1 - page 19, line 6; Figures 5-9, all elements).

B. CLAIM 12 - INDEPENDENT

Specifically with respect to Claim 12, such claim is directed to a computer program product in a computer readable medium for multicasting a data packet in a system area network. First instructions are provided for receiving the data packet, where the data packet includes an identifier of a multicast group. Second instructions are provided for identifying a plurality of queue pairs that are members of the multicast group. Third instructions are provided for delivering the data packet to each of the plurality of queue pairs that are members of the multicast group (Specification page 9, line 14 – page 10, line 2; Figure 3, element 300; Specification page 11, line 1 - page 19, line 6; Figures 5-9, all elements).

C. CLAIM 23 – INDEPENDENT and MEANS-PLUS-FUNCTION

Specifically with respect to Claim 23, such claim is directed to an apparatus for multicasting a data packet in a system area network. The apparatus includes means for receiving a data packet, where the data packet includes an identifier of a multicast group. The apparatus also includes means for identifying a plurality of queue pairs that are members of the multicast group. The apparatus also includes means for delivering the data packet to each of the plurality

of queue pairs that are members of the multicast group (Specification page 9, line 14 – page 10, line 2; Figure 3, element 300; Specification page 11, line 1 - page 19, line 6; Figures 5-9, all elements).

The corresponding structure for each of the three means-for elements ('means for receiving', 'means for identifying' and 'means for delivering') are depicted in Figure 3, element 300 and described at Specification page 9, line 14 – page 10, line 2.

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

The grounds of rejection to review on appeal are as follows:

- A.** Whether Claims 1-35 are anticipated by Kashyap (US 6,944,786) under 35 U.S.C. § 102(e).

- B.** Whether the Examiner improperly failed to enter or consider the affidavit or other evidence submitted by Appellants in their Response to Final Office Action submission dated July 3, 2006.

ARGUMENT

A. GROUND OF REJECTION 1 (Claims 1-35)

The Examiner rejected Claims 1-35 under 35 U.S.C. § 102(e) as being anticipated by Kashyap (US 6,944,786). Appellants show clear error in such rejection as follows.

For a prior art reference to anticipate in terms of 35 U.S.C. 102, every element of the claimed invention must be identically shown in a single reference. *In re Bond*, 910 F.2d 831, 15 USPQ2d 1566 (Fed. Cir. 1990) (emphasis added by Appellants). Appellants will now show that every element of the claimed invention is not identically shown in a single reference, and thus Claims 1-35 are not anticipated by the cited reference, and therefore the rejection of such claims is clearly erroneous.

A.1. Claims 1, 2, 4-6, 12, 13, 15-17, 23, 24, 26-28 and 34

With respect to Claim 1, it is urged that the cited reference does not teach the claimed features of (1) “identifying *a plurality of queue pairs* that are members of the multicast group” or (2) “delivering the data packet to *each of the plurality of queue pairs* that are members of the multicast group”. Instead, Kashyap’s discussion pertaining to multicast groups states that *ports* are associated with a multicast group (Kashyap col. 6, lines 44-56) – and not *queue pairs* as claimed.

In rejecting the ‘identifying’ step of Claim 1, the Examiner cites Kashyap’s teachings at col. 6, line 10; col. 5, line 52 as teaching ‘identifying a plurality of queue pairs’ and Kashyap’s teachings at col. 5, lines 51-56 as teaching ‘that are members of the multicast group’. Appellants urge that Kashyap states at cited col. 6, line 10:

“The end node 502 has running thereon processes 504 having associated QPs 506, 508, and 510.”

and states at cited col. 5, line 52:

“Each process may have associated therewith one or more queue pairs (QPs)”

and states at cited col. 5, lines 51-56:

“Each process may have associated therewith one or more queue pairs (QPs), where each QP communicates with the channel adapter (CA) 418 of the node 400 to link to the Infiniband fabric, as indicated by the arrow 420. For example, the process 402 specifically has QPs 406 and 408, whereas the process 404 has a QP 410.”

As can be seen, none of these passages describe any type of multicast group, or of identifying *a plurality of queue pairs that are members of a multicast group*. Rather, these cited passages discuss associating queue pairs with *processes executing within a node* (e.g., see Figure 4 where process A has associated queue pairs 406 and 408 and process B has an associated queue pair 410).

As to Kashyap’s processing of multicast group information (which is described elsewhere, and is not described in these cited passages), such multicast processing is with respect to a *given port of a channel adapter*, and is not related to particular queue pairs (Kashyap col. 6, lines 46-48 – “In general, a *set of end nodes may join a multicast group*, such that the SM *assigns a port of each node* with a multicast DLID of the multicast group”; Kashyap col. 6, lines 62-64 – “Data packets received by the switch 504 that specify the multicast DLID are thus sent from one of these multicast ports *to the associated ports of the multicast group nodes*”). A review of Kashyap’s Figure 5 clearly shows that ports (elements 514, 516 and 518) are very different from queue pairs (elements 506, 508 and 510). Thus, Kashyap’s association of a multicast DLID with a *port of each node* that is a member of the multicast group does not teach any step of associating or identifying *queue pairs that are members of a multicast group*.

Per the invention of Claim 1, a data packet is received, and this received data packet includes an identifier of a multicast group, and it is this multicast group for which *a plurality of queue pairs are identified that are members of such multicast group*. The Kashyap processes, which are associated with queue pairs, are *not identified by a received data packet*, and thus it is error to equate Kashyap’s processes with the claimed multicast group. The Kashyap ports, to which a multicast data packet is associated and sent, are very different from queue pairs. The Kashyap multicast group is directed to and associated with a *port* of each node (Kashyap Figure

5, element 514, 516), and not to *queue pairs*. Kashyap teaches that a multicast data packet is sent to ports of the multicast group nodes. There is no association between a multicast group and queue pairs. In fact, Kashyap teaches that the operation of each queue pair is independent from the others (col. 6, line 1), further evidencing that multiple queue pairs are not grouped/associated with a given multicast group.

In summary, per Kashyap's discussion regarding a received packet having a multicast DLID, such packet is sent to a *particular port for each node that has joined the group* (Kashyap col. 3, lines 22-56; col. 6, lines 44-65 and particularly lines 56-65). The present invention advantageously allows for identifying *specific queue pairs that are members of a multicast group* such that a received data packet can be delivered to each of such queue pairs. Quite simply, the cited reference associates nodes (and a given port of each node) with a multicast group, whereas Claim 1 identifies *specific queue pairs* that are members of a multi-cast group. Thus, it is urged that Claim 1 is clearly not anticipated by the cited reference, as every element of the claimed invention is not identically shown in a single reference.

In rejecting the 'delivering' step of Claim 1, the Examiner cites Kashyap's teaching at col. 5, lines 59-64. Appellants urge that Kashyap states at this cited passage:

"A QP includes a send work queue and a receive work queue that are paired together. In general, the send work queue holds instructions that cause data to be transferred between the client's memory and another process's memory, and the receive work queue holds instructions about where to place data that is received from another process."

As can be seen, this passage describes particulars of a single queue pair that holds instructions. In contrast, Claim 1 specifically recites an active step of *delivering* a data packet, where the data packet is delivered to each of the plurality of queue pairs that are members of the multicast group. This cited passage provides no such teaching, and in fact makes no mention of any type of data packet delivery to members of a multicast group, either as claimed or otherwise, and thus it is further urged that Claim 1 is not anticipated by the cited reference as there is yet another claimed feature not identically shown in a single reference, *In re Bond, Id.*

A.2 Claims 3, 14 and 25

Appellants initially urge clear error in the rejection of Claim 3 (and similarly for Claims 14 and 25) for reasons given above with respect to Claim 1 (of which Claim 3 depends upon). Still further, Claim 3 recites “wherein the data packet is received in a channel adapter of an end node, wherein the data packet originates from the end node, and wherein delivering the data packet to each of the plurality of queue pairs that are members of the multicast group includes replicating the data packet for each of the plurality of queue pairs that are internal to the end node”. As can be seen, per the features of Claim 3, the end node that receives the data packet is the same end node that initiated the data packet, and the delivering step of Claim 1 is further refined to require that delivering the data packet to each of the plurality of queue pairs that are members of the multicast group includes replicating the data packet for each of the plurality of queue pairs that are internal to the end node (which originated the data packet). The cited reference does not teach the replication of data packets for each of a plurality of queue pairs that are internal to an end node which both originates and receives a given data packet.

In rejecting Claim 3, the Examiner states that this claimed step is inherent in the Kashyap teachings. Appellants urge error in such inherency assertion. “To establish inherency,” the Federal Circuit recently stated, “the extrinsic evidence `must make clear that the missing descriptive matter is necessarily present in the thing described in the reference, and that it would be so recognized by persons of ordinary skill.” *In re Robertson*, 169 F.3d 743, 745 [49 USPQ2d 1949] (Fed. Cir. 1999); see also *Continental Can Co. U.S.A., Inc. v. Monsanto Co.*, 948 F.2d 1264, 1268 [20 USPQ2d 1746] (Fed. Cir. 1991). Such inherency may not be established by “probabilities or possibilities.” *Continental Can*, 948 F.2d at 1269 (quoting *In re Oelrich*, 666 F.2d 578, 581 [212 USPQ 323] (C.C.P.A. 1981)). Because Kashyap does not associate queue pairs with a broadcast group – but instead associates ports with a broadcast group – there would be no reason to replicate data packets *for each of a plurality of queue pairs*. At best, Kashyap describes sending a data packet that specifies a multicast group to a plurality of *ports* of the multicast group (Kashyap col. 6, lines 62-65). Thus, as every element of the claimed invention is not identically shown in a single reference, it is further shown that Claim 3 has been erroneously rejected under 35 USC § 102(e).

A.3 Claims 7, 18 and 29

Appellants initially urge clear error in the rejection of Claim 7 (and similarly for Claims 18 and 29) for reasons given above with respect to Claim 1 (of which Claim 7 depends upon). Still further, Claim 7 recites “wherein identifying the plurality of queue pairs includes determining which queue pairs are associated with a destination local identifier in the data packet”. In rejecting this aspect of Claim 7, the Examiner cites Kashyap’s discussion at col. 6, lines 56-66 as teaching this claimed feature. Appellants urge that there, Kashyap states:

“When a data packet that has a multicast DLID is received, the multicast DLID is examined, and the data packet is forwarded, based on the tables programmed by the SM. If the multicast DLID is not in the table, or the switch does not maintain tables, that it forwards the packets on the primary and non-primary default multicast **ports**. Data packets received by the switch 504 that specify the multicast DLID are thus sent from one of these multicast ports to the **associated ports of the multicast group nodes**. The switch 504 can be configured with routing information for the multicast traffic that specifies the **ports** where the packet should travel” (emphasis added by Appellants)

This cited passage makes no mention whatsoever of any type of queue pair or queue pair identification, and thus cannot teach the specific claimed step of “wherein identifying the *plurality of queue pairs includes determining which queue pairs are associated with a destination local identifier in the data packet*” (emphasis added by Appellants). As shown above with respect to Claim 1, a port is very different from a queue pair, and a description of identifying *ports* for a received multicast data packet does not teach identifying *queue pairs* that are associated with either a data packet or a destination local identifier of such a data packet. Thus, as every element of the claimed invention is not identically shown in a single reference, it is further shown that Claim 7 has been erroneously rejected under 35 USC § 102(e).

A.4 Claims 8, 19 and 30

Appellants initially urge clear error in the rejection of Claim 8 (and similarly for Claims 19 and 30) for reasons given above with respect to Claim 7 (of which Claim 8 depends upon). Still further, Claim 8 recites “wherein determining which queue pairs are associated with the destination local identifier includes using a destination local identifier to queue pair lookup table”. As can be seen, the queue pair determination is done using a destination local identifier to queue pair lookup table. In rejecting this aspect of Claim 8, the Examiner cites Kashyap’s discussion at col. 6, lines 56-66 as teaching this claimed feature. Appellants urge that there, Kashyap states:

“When a data packet that has a multicast DLID is received, the multicast DLID is examined, and the data packet is forwarded, based on the tables programmed by the SM. If the multicast DLID is not in the table, or the switch does not maintain tables, that it forwards the packets on the primary and non-primary default multicast **ports**. Data packets received by the switch 504 that specify the multicast DLID are thus sent from one of these multicast ports to the **associated ports of the multicast group nodes**. The switch 504 can be configured with routing information for the multicast traffic that specifies the **ports** where the packet should travel” (emphasis added by Appellants)

This cited passage makes no mention whatsoever of any type of queue pair or queue pair table, and thus cannot teach the specific claimed feature of “wherein determining which queue pairs are associated with the destination local identifier includes using a destination local identifier to queue pair lookup table” (emphasis added by Appellants). Thus, as every element of the claimed invention is not identically shown in a single reference, it is further shown that Claim 8 has been erroneously rejected under 35 USC § 102(e).

A.5 Claims 9, 20 and 31

Appellants initially urge clear error in the rejection of Claim 9 (and similarly for Claims 20 and 31) for reasons given above with respect to Claim 8 (of which Claim 9 depends upon).

Still further, Claim 9 recites “wherein the destination local identifier to queue pair lookup table contains a fixed number of queue pair identifier columns per destination local identifier”. As can be seen, the destination local identifier to queue pair lookup table contains a fixed number of queue pair identifier columns per destination local identifier. In rejecting this aspect of Claim 9, the Examiner cites Kashyap’s discussion at col. 6, lines 56-66 as teaching this claimed feature. Appellants urge that there, Kashyap states:

“When a data packet that has a multicast DLID is received, the multicast DLID is examined, and the data packet is forwarded, based on the tables programmed by the SM. If the multicast DLID is not in the table, or the switch does not maintain tables, that it forwards the packets on the primary and non-primary default multicast **ports**. Data packets received by the switch 504 that specify the multicast DLID are thus sent from one of these multicast ports to the **associated ports of the multicast group nodes**. The switch 504 can be configured with routing information for the multicast traffic that specifies the **ports** where the packet should travel” (emphasis added by Appellants)

This cited passage makes no mention whatsoever of any type of queue pair or queue pair table, and thus cannot teach the specific claimed feature of “wherein the destination local identifier to queue pair lookup table contains a *fixed number of queue pair identifier columns per destination local identifier*” (emphasis added by Appellants). Thus, as every element of the claimed invention is not identically shown in a single reference, it is further shown that Claim 9 has been erroneously rejected under 35 USC § 102(e).

A.6 Claims 10, 21 and 32

Appellants initially urge clear error in the rejection of Claim 10 (and similarly for Claims 21 and 32) for reasons given above with respect to Claim 8 (of which Claim 10 depends upon). Still further, Claim 10 recites “wherein the destination local identifier to queue pair lookup table contains a flexible number of queue pair identifier columns per destination local identifier”. As can be seen, the destination local identifier to queue pair lookup table contains a flexible

number of queue pair identifier columns per destination local identifier. In rejecting this aspect of Claim 10, the Examiner cites Kashyap's discussion at col. 6, lines 56-66 as teaching this claimed feature. Appellants urge that there, Kashyap states:

“When a data packet that has a multicast DLID is received, the multicast DLID is examined, and the data packet is forwarded, based on the tables programmed by the SM. If the multicast DLID is not in the table, or the switch does not maintain tables, that it forwards the packets on the primary and non-primary default multicast **ports**. Data packets received by the switch 504 that specify the multicast DLID are thus sent from one of these multicast ports to the **associated ports of the multicast group nodes**. The switch 504 can be configured with routing information for the multicast traffic that specifies the **ports** where the packet should travel” (emphasis added by Appellants)

This cited passage makes no mention whatsoever of any type of queue pair or queue pair table, and thus cannot teach the specific claimed feature of “wherein the destination local identifier to queue pair lookup table contains a *flexible number of queue pair identifier columns* per destination local identifier” (emphasis added by Appellants). Thus, as every element of the claimed invention is not identically shown in a single reference, it is further shown that Claim 10 has been erroneously rejected under 35 USC § 102(e).

A.7 Claims 11, 22 and 33

Appellants initially urge clear error in the rejection of Claim 11 (and similarly for Claims 22 and 33) for reasons given above with respect to Claim 10 (of which Claim 11 depends upon). Still further, Claim 11 recites “wherein one of the queue pair identifier columns associated with the destination local identifier *serves as a link to another entry* in the destination local identifier to queue pair lookup table” (emphasis added). As can be seen, one of the queue pair identifier columns (that is associated with the destination local identifier) serves as a link to another entry in the destination local identifier to queue pair lookup table. In rejecting this aspect of Claim 11, the Examiner cites Kashyap's discussion at col. 6, lines 56-66 as teaching this claimed feature.

Appellants urge that there, Kashyap states:

“When a data packet that has a multicast DLID is received, the multicast DLID is examined, and the data packet is forwarded, based on the tables programmed by the SM. If the multicast DLID is not in the table, or the switch does not maintain tables, that it forwards the packets on the primary and non-primary default multicast **ports**. Data packets received by the switch 504 that specify the multicast DLID are thus sent from one of these multicast ports to the **associated ports of the multicast group nodes**. The switch 504 can be configured with routing information for the multicast traffic that specifies the **ports** where the packet should travel” (emphasis added by Appellants)

This cited passage makes no mention whatsoever of any type of queue pair or queue pair table, and thus cannot teach the specific claimed feature of “wherein *one of the queue pair identifier columns* associated with the destination local identifier *serves as a link to another entry* in the destination local identifier to *queue pair lookup table*” (emphasis added by Appellants). Thus, as every element of the claimed invention is not identically shown in a single reference, it is further shown that Claim 11 has been erroneously rejected under 35 USC § 102(e).

A.8 Claim 35

Appellants initially urge clear error in the rejection of Claim 35 for reasons given above with respect to Claim 1 (of which Claim 35 depends upon). Still further, Claim 35 recites “wherein delivering the data packet to each of the plurality of queue pairs that are members of the multicast group includes: determining if there is an error in delivering the data packet to each of the plurality of queue pairs; and dropping the data packet if an error occurs during delivery of the data packet to each of the plurality of queue pairs”. As can be seen, such claim is directed to (i) determining if there is an error in data packet delivery to each of the plurality of queue pairs, and (ii) if there is such an error, the data packet is dropped – both of which are substeps of the ‘delivering’ step of Claim 1. In rejecting Claim 34, the Examiner cites Kashyap’s teaching at col. 5, lines 1-10 as teaching every element recited in Claim 35. Applicants urge that there,

Kashyap states:

“Furthermore, InfiniBand networks allow for a number of different transport services, including reliable and unreliable connections, reliable and unreliable datagrams, and raw packet support. In reliable connections and datagrams, acknowledgments and packet sequence numbers for guaranteed packet ordering are generated. Duplicate packets are rejected, and missing packets are detected. In unreliable connections and datagrams, acknowledgments are not generated, and packet ordering is not guaranteed. Duplicate packets may not be rejected, and missing packets may not be detected.”

As can be seen, this passage describes reliable connections and datagrams (where acknowledgements and packet sequence numbers are generated), and unreliable connections and datagrams (where acknowledgements are not generated and packet ordering is not guaranteed). There is also mention of duplicate packets may or may not be rejected, and missing packets may or may not be detected. There is no teaching or discussion of any type pertaining to queue pairs, and thus this cited passage cannot teach any type of determination as to determining if there is an error in *delivering the data packet to each of the plurality of queue pairs*. Thus, as every element of the claimed invention is not identically shown in a single reference, it is further shown that Claim 35 has been erroneously rejected under 35 USC § 102(e).

B. Failure to Enter or Consider Affidavit or Other Evidence

In an Advisory Action dated 08/02/06, the Examiner refused to enter the affidavit or other evidence submitted by Appellants in their Response to Final Office Action submission dated July 3, 2006. In denying such entry, the Examiner stated that applicant failed to provide a showing of good and sufficient reasons why the affidavit or other evidence is necessary and was not presented earlier.

Appellants respectfully point out that this affidavit and other evidence was provided by Appellants as a direct result of an expressed suggestion and request for such information by the Examiner. In an Office Action dated December 2, 2005, on page 2 of such action, the Examiner stated that the 35 USC 102 rejection might be overcome “by an appropriate showing under 37

CFR 1.131". Instead of continuing to debate the merits of patentability with respect to the cited reference, Appellants instead chose to follow the path of least resistance and provide an appropriate showing under 37 CFR 1.131 – as suggested by the Examiner – such that this case could expeditiously pass to issuance. It is urged that the Examiner's suggestion/solicitation of a 37 CFR 1.131 affidavit to overcome the 35 USC 102 rejection now requires that the Examiner enter and consider such affidavit and other evidence, as such was submitted as a direct result/consequence of such suggestion/solicitation. Appellants justifiably relied upon the Examiner's suggested course of action – and in reliance thereof chose to not further argue or debate the merits of the cited reference, but instead to provide the suggested affidavit and other evidence – in order to expeditiously further the prosecution of this case and to make a bona fide attempt to advance the application to a condition for allowance. It is clear error for the Examiner to request/suggest such 37 CFR 1.131 showing, and then refuse entry/consideration of same.

In conclusion, Appellants have shown numerous and substantial error in the final rejection of all pending claims in this case, and respectfully requests that the Board reverse such rejection.

/Wayne P. Bailey/
Wayne P. Bailey
Reg. No. 34,289
YEE & ASSOCIATES, P.C.
PO Box 802333
Dallas, TX 75380
(972) 385-8777

CLAIMS APPENDIX

The text of the claims involved in the appeal are:

1. A method of multicasting a data packet in a system area network, comprising:
receiving the data packet, wherein the data packet includes an identifier of a multicast group;
identifying a plurality of queue pairs that are members of the multicast group; and
delivering the data packet to each of the plurality of queue pairs that are members of the multicast group.
2. The method of claim 1, wherein the data packet is received in a channel adapter of an end node, the end node being a final destination for the data packet within the system area network.
3. The method of claim 1, wherein the data packet is received in a channel adapter of an end node, wherein the data packet originates from the end node, and wherein delivering the data packet to each of the plurality of queue pairs that are members of the multicast group includes replicating the data packet for each of the plurality of queue pairs that are internal to the end node.
4. The method of claim 1, further comprising:
decoding the data packet; and
storing the data packet in a multicast packet buffer.

5. The method of claim 4, wherein decoding the data packet and storing the data packet in the multicast packet buffer are performed by port logic.
6. The method of claim 4, wherein decoding the data packet and storing the data packet in the multicast packet buffer are performed by channel adapter logic.
7. The method of claim 1, wherein each of the plurality of queue pairs comprises a send queue and a receive queue, and wherein identifying the plurality of queue pairs includes determining which queue pairs are associated with a destination local identifier in the data packet.
8. The method of claim 7, wherein determining which queue pairs are associated with the destination local identifier includes using a destination local identifier to queue pair lookup table.
9. The method of claim 8, wherein the destination local identifier to queue pair lookup table contains a fixed number of queue pair identifier columns per destination local identifier.
10. The method of claim 8, wherein the destination local identifier to queue pair lookup table contains a flexible number of queue pair identifier columns per destination local identifier.
11. The method of claim 10, wherein one of the queue pair identifier columns associated with the destination local identifier serves as a link to another entry in the destination local identifier to queue pair lookup table.

12. A computer program product in a computer readable medium for multicasting a data packet in a system area network, comprising:

first instructions for receiving the data packet, wherein the data packet includes an identifier of a multicast group;

second instructions for identifying a plurality of queue pairs that are members of the multicast group; and

third instructions for delivering the data packet to each of the plurality of queue pairs that are members of the multicast group.

13. The computer program product of claim 12, wherein the data packet is received in a channel adapter of an end node, the end node being a final destination for the data packet within the system area network.

14. The computer program product of claim 12, wherein the data packet is received in a channel adapter of an end node, wherein the data packet originates from the end node, and wherein the third instructions for delivering the data packet to each of the plurality of queue pairs that are members of the multicast group include instructions for replicating the data packet for each of the plurality of queue pairs that are internal to the end node.

15. The computer program product of claim 12, further comprising:

fourth instructions for decoding the data packet; and

fifth instructions for storing the data packet in a multicast packet buffer.

16. The computer program product of claim 15, wherein the fourth instructions for decoding the data packet and the fifth instructions for storing the data packet in the multicast packet buffer are executed by port logic.

17. The computer program product of claim 15, wherein the fourth instructions for decoding the data packet and the fifth instructions for storing the data packet in the multicast packet buffer are executed by channel adapter logic.

18. The computer program product of claim 12, wherein each of the plurality of queue pairs comprises a send queue and a receive queue, and wherein the second instructions for identifying the plurality of queue pairs include instructions for determining which queue pairs are associated with a destination local identifier in the data packet.

19. The computer program product of claim 18, wherein the instructions for determining which queue pairs are associated with the destination local identifier include instructions for using a destination local identifier to queue pair lookup table.

20. The computer program product of claim 19, wherein the destination local identifier to queue pair lookup table contains a fixed number of queue pair identifier columns per destination local identifier.

21. The computer program product of claim 19, wherein the destination local identifier to queue pair lookup table contains a flexible number of queue pair identifier columns per destination local identifier.

22. The computer program product of claim 21, wherein one of the queue pair identifier columns associated with the destination local identifier serves as a link to another entry in the destination local identifier to queue pair lookup table.

23. An apparatus for multicasting a data packet in a system area network, comprising:
means for receiving the data packet, wherein the data packet includes an identifier of a multicast group;
means for identifying a plurality of queue pairs that are members of the multicast group;
and
means for delivering the data packet to each of the plurality of queue pairs that are members of the multicast group.

24. The apparatus of claim 23, wherein the data packet is received in a channel adapter of an end node, the end node being a final destination for the data packet within the system area network.

25. The apparatus of claim 23, wherein the data packet is received in a channel adapter of an end node, wherein the data packet originates from the end node, and wherein the means for delivering the data packet to each of the plurality of queue pairs that are members of the

multicast group includes means for replicating the data packet for each of the plurality of queue pairs that are internal to the end node.

26. The apparatus of claim 23, further comprising:

means for decoding the data packet; and

means for storing the data packet in a multicast packet buffer.

27. The apparatus of claim 26, wherein the means for decoding the data packet and means for storing the data packet in the multicast packet buffer are include port logic.

28. The apparatus of claim 26, wherein the means for decoding the data packet and means for storing the data packet in the multicast packet buffer include channel adapter logic.

29. The apparatus of claim 23, wherein each of the plurality of queue pairs comprises a send queue and a receive queue, and wherein the means for identifying the plurality of queue pairs includes means for determining which queue pairs are associated with a destination local identifier in the data packet.

30. The apparatus of claim 29, wherein the means for determining which queue pairs are associated with the destination local identifier includes means for using a destination local identifier to queue pair lookup table.

31. The apparatus of claim 30, wherein the destination local identifier to queue pair lookup table contains a fixed number of queue pair identifier columns per destination local identifier.
32. The apparatus of claim 30, wherein the destination local identifier to queue pair lookup table contains a flexible number of queue pair identifier columns per destination local identifier.
33. The apparatus of claim 32, wherein one of the queue pair identifier columns associated with the destination local identifier serves as a link to another entry in the destination local identifier to queue pair lookup table.
34. The method of claim 1, wherein receiving the data packet includes:
determining if there is an error in receipt of the data packet; and
if there is an error in receipt of the data packet, dropping the data packet.
35. The method of claim 1, wherein delivering the data packet to each of the plurality of queue pairs that are members of the multicast group includes:
determining if there is an error in delivering the data packet to each of the plurality of queue pairs; and
dropping the data packet if an error occurs during delivery of the data packet to each of the plurality of queue pairs.

EVIDENCE APPENDIX

There is no evidence to be presented.

RELATED PROCEEDINGS APPENDIX

There are no related proceedings.